

Single Constant Multiplication for SAT

Hendrik Bierlee^{1,2}, Jip J. Dekker^{1,2}, Vitaly Lagoon³, Peter J. Stuckey^{1,2}, and Guido Tack^{1,2}

¹ Monash University, Wellington Rd, Clayton VIC 3800, Australia
{hendrik.bierlee,jip.dekker,peter.stuckey,guido.tack}@monash.edu

² OPTIMA ARC Industrial Training and Transformation Centre, Australia

³ Cadence Design Systems, USA
lagoon@cadence.com

Abstract. This paper presents new methods of encoding the multiplication of a binary encoded integer variable with a constant value for Boolean Satisfiability (SAT) solvers. This problem, known as the Single Constant Multiplication (SCM) problem, is widely studied for its application in hardware design, but this technique are currently not employed for SAT encodings. Considering the smaller and variable bit sizes and the different cost of operations in SAT, we propose further improvements to existing methods by minimising the number of full/half adders, rather than the number of ripple carry adders. We compare our methods to simple encodings into adders, currently employed by SAT encoders, and direct Boolean encoding using logic minimisation tools. We show that our methods result in improved solve-time for problems involving integer linear constraints. A library of optimal recipes for each method to encode SCM for SAT is made available as part of this publication.

Keywords: Single Constant Multiplier · Boolean Satisfiability · Encoding constraints

1 Introduction

Boolean Satisfiability (SAT) is a powerful approach to solving combinatorial problems, but in order use a SAT solver, we need to encode the problem into clauses. One fundamental constraint is the constant multiplication $y = c \times x$ of input and output integer variables x and y , and constant c . If x and y are encoded using a binary representation, any such multiplication can be achieved using a combination of additions, subtractions, and multiplications by a power of two. For example, the multiplication $5 \times x$ can be decomposed into $4 \times x + 1 \times x$. Since multiplication by a power of two in a binary encoding is a simple left-shift operation, which does not incur any cost, this decomposition effectively turns the multiplication into a sequence of additions.

The most basic decomposition for a multiplication $c \times x$ therefore simply considers the binary representation of c and introduces one addend for every 1 in this binary representation – such as in the example of $5x = 1x + 4x$ above, or $117x = 64x + 32x + 16x + 4x + 1x$. However, decompositions that reuse

intermediate results or use a combination of additions and subtractions can result in fewer additions. For example, we could define $63x = 64x - x$, $59x = 63x - 4x$, $117x = 2 \times 59x - x$. This only requires 3 additions/subtractions, compared to the 4 in the original example.

In hardware applications such as ASIC and FPGA, it is important for the performance and cost of the application to use small circuits, and therefore there is considerable previous research on finding a circuit for a given constant that minimizes the number k of additions and subtractions. This is known as the *Single Constant Multiplication* (SCM) problem. SCM is NP-hard [7], and has an upper bound of $k \leq \lceil \log_2(c)/2 \rceil$ [3]. For many practical instances of the problem, however, k is known to be much smaller.

When encoding constant multiplication for SAT, an approach similar to SCM can be used since integer variables – like in hardware – are often encoded in a binary representation [5]. Unlike hardware, the number of bits w to represent x is not fixed, but depends on the domain of x (e.g., only $w = 2$ bits are required to represent $x \in [0, 2]$ with Boolean variables x_0, x_1).

Furthermore, the complexity of the SAT encoding of a single addition or subtraction (using a ripple carry adder [17]) varies. The required number of full (and half) adders is determined by the constant as well as the number of bits of the inputs. In fact, the addition of $x + 4x = 5x$ requires *no* adders if x has two bits x_0, x_1 , since $5x$ would be represented using the existing bits x_0, x_1, x_0, x_1 . If x has three bits, multiple adders are required. Consequently, instead of minimizing the number of additions/subtractions k , a better circuit for SAT minimizes the total number of full/half adders a .

Example 1. Consider two circuits for $c = 117$: $17x = 16x + x$; $19x = 2x + 17x$; $117x = 8 \times 17x (= 136x) - 19x$ and $63x = 64x - x$; $59x = 63x - 4x$; $117x = 2 \times 59x (= 118x) - x$. Both are minimal with respect to $k = 3$, but assuming $w = 4$ bits for $x \in [0, 7]$, the first requires $a = 18$ full/half adders, while the second requires $a = 31$. Yet another circuit, $3x = 2x + x$; $49x = 16 \times 3x (= 48x) + x$; $17x = 16x + x$; $117x = 4 \times 17x (= 68x) + 49x$, requires only $a = 12$ full/half adders, even though it uses $k = 4$ additions/subtractions. For $w = 4$, $a = 12$ is optimal for $c = 117$, but for different w the best circuit changes.

Note that if we constrain all intermediates to be represented in a fixed bit width (common in hardware solutions), then this may change the answer. For example, given fixed 12 bit width for intermediates the first recipe is not applicable when x is a 5 bit number, since intermediate $136x$ is not guaranteed to fit in 12 bits, while $117x$ does. $\square$

The main contribution of this paper is the application of the SCM methodology in the context of SAT, the extension to the basic approach that result in better SAT encodings by minimizing full/half adders, and a database with pre-computed *recipes* of the SCM problem for a range of constants and bit widths. Our database exhaustively covers the parameter space over $1 \leq w \leq 16$ and $1 \leq c \leq 2047$. We assume no fixed bit width restriction on intermediate results.

We present a set of recipes to compute $c \times x$ optimal in terms of number of additions/subtractions for the general SCM problem, independent of the size of x , assuming we can use any number of bits for intermediate results. Then, we produce a set of recipes that minimize the number of full/half adders to compute $c \times x$. These recipes differ depending on the bit size of x . The SCM SAT encoding database is available as part of the released implementation and benchmarks [8].

We evaluate the different SCM recipes in section 4, where we use the pre-computed circuits to encode problems which focus on constant multiplication. We observe significant improvements in terms of the size of the final encoding and the SAT solver performance. In section 5, we discuss future work and conclude our findings.

2 Preliminaries and Related Work

In this section, we give preliminaries on the *Constraint Satisfaction Problem* (CSP) and SAT problem. We then formally define the SCM problem and discuss existing approaches for solving it.

2.1 Constraint Satisfaction Problem

A CSP [15] instance, $P = (\mathcal{X}, D, \mathcal{C})$, consists of a set of variables $\mathcal{X}$, with each $x \in \mathcal{X}$ restricted to take values from some domain $D(x)$. For this paper we assume domains are ordered sets of integers, and denote by $\text{lb}(x)$ and $\text{ub}(x)$ the least and greatest values in $D(x)$. We will use interval notation $[l, u]$ to represent the set of integers $\{l, l+1, \dots, u\}$. A set of constraints $\mathcal{C}$ expresses relationships between the variables. An *assignment* of a CSP instance is a mapping of variables to values. An assignment that is consistent with the domains and constraints is a *solution* to the instance. A *Constraint Optimization Problem* (COP) is a CSP with an additional objective function f that assigns a value to each assignment. An optimal solution to a minimization (maximization) COP is one that minimizes (maximizes) f .

2.2 Boolean Satisfiability

A SAT problem can be seen as a special case of a CSP, where the domain for all variables x is $D(x) \in \{0, 1\}$, representing the values *false* and *true*. A *literal* is either a Boolean variable x or its negation $\neg x$. We extend the negation operation to operate on literals, i.e., $\neg l = \neg x$ if $l = x$ and $\neg l = x$ if $l = \neg x$. We use the notation $l = v$ where l is a literal and $v \in \{0, 1\}$ to encode the appropriate form of the literal, i.e., if $v = 1$ it is equivalent to l and if $v = 0$ it is equivalent to $\neg l$. A *clause* is a disjunction of literals. In a SAT problem P , the constraints are in *Conjunctive Normal Form* (CNF), which means that $\mathcal{C}$ is a set of clauses.

A general CSP can be mapped, or encoded, into SAT, by encoding each integer variable into a set of Boolean *encoding* variables, and each constraint into a set of clauses and additional auxiliary Boolean variables. There are a

number of different ways of encoding integer variables [5], such as the direct and order encodings. In this paper, we focus on the *binary* encoding.

Given an integer variable x with non-negative, possibly non-contiguous domain $D(x)$, the unsigned binary encoding method maps x to $m_x = \lfloor \log_2 \text{ub}(x) \rfloor + 1$ encoding variables $\llbracket \text{bit}(x, i) \rrbracket, 0 \leq i < m_x$. The semantics of the encoding can be expressed by the equation $x = \sum_{i=0}^{m_x-1} 2^i \times \llbracket \text{bit}(x, i) \rrbracket$. For example, for integer variables d with domain $D(d) = [0, 127]$, an encoding with $m_d = 7$ bits can be used. In an assignment with $d = 117 = 0b1110101$, we would have $\llbracket \text{bit}(d, 0) \rrbracket = 1$, $\llbracket \text{bit}(d, 1) \rrbracket = 0$, $\llbracket \text{bit}(d, 2) \rrbracket = 1$ and so on. If the initial domain $D(x)$ is non-contiguous, or has bounds that are not powers of 2, these have to be enforced using additional constraints (clauses). The representation can be extended in a straightforward way to support negative domain values, by assuming a two's complement encoding instead of a simple unsigned binary encoding. For the rest of the paper, we will assume non-negative domains to simplify presentation, and only discuss the general approach where required.

2.3 The Single Constant Multiplication Problem

The SCM problem can be formalized as follows. Given the multiplier target c , we are to decide a sequence of up to n equations of the form $c' \times x = 2^{s_l} \times (c_l \times x) \pm 2^{s_r} \times (c_r \times x)$. The left argument $c_l \times x$ is either a previously computed multiple or $c_l = 1$, and $s_l \geq 0$ is the left-shift of the left argument. Similarly, $c_r \times x$ is the right argument, and $s_r \geq 0$ is its left-shift.⁴ The final equation should define $c \times x$.

We can thus encode a recipe for $c = c^k$ as a sequence of k equations represented by tuples $\langle c_l^i, s_l^i, c_r^i, s_r^i, \#^i, c^i \rangle$, where $c^i = (c_l^i \ll s_l^i) \#^i (c_r^i \ll s_r^i)$ such that $k \leq n$, $\#^i \in \{+, -\}$, $\{c_l^i, c_r^i\} \subseteq \{c^j \mid j < i\} \cup \{1\}$ and $s_l^i, s_r^i \geq 0$.

Example 2. The first recipe for 117 of example 1 is defined by the sequence $[\langle 1, 4, 1, 0, +, 17 \rangle, \langle 1, 1, 17, 0, +, 19 \rangle, \langle 17, 3, 19, 0, -, 117 \rangle]$. The second recipe is defined by $[\langle 1, 6, 1, 0, -, 63 \rangle, \langle 63, 0, 1, 2, -, 59 \rangle, \langle 59, 1, 1, 0, -, 117 \rangle]$.

Usually, the goal of the SCM problem is to find the minimal k recipe for a given c (knowing that it is bounded by $k \leq \lceil \log_2(c)/2 \rceil$ [3]).

Related work The SCM problem has been studied extensively in the context of hardware. Many approaches for finding optimal circuits have been proposed, using graph based techniques [9], or using *Integer Linear Programming* (ILP) solvers [11], SAT solvers [12,2], as well as heuristic approaches [2]. A generalization of SCM is the MCM, which builds a circuit for multiple target constants. Compared to SCM, this enables further sharing of intermediate results between the circuits for different constants. Apart from minimizing the number of nodes

⁴ Some versions of SCM also allows right shifts, and while these are necessary for optimal *Multiple Constant Multiplication* (MCM) solutions, we did not come across any instances of SCM where they improved the solution.

(additions and subtractions), some of these related works consider other objectives such as minimizing the required surface area, delay, or power consumption of the generated circuits. These metrics are obviously useful when generating hardware circuits, but they are not directly relevant for SAT encodings. We are not aware of approaches that minimize the number of full/half adders in the resulting circuit. This metric is useful in the context of SAT solvers, since it results in a smaller CNF in terms of the number of variables, clauses and literals, as we shall see in section 4. We briefly discuss MCM for SAT as future work in section 5.

Another difference with existing approaches is that the domains of intermediate results is left unbounded when solving SCM for SAT. Approaches that model hardware circuits have a fixed bit width h for all results $c' \times x$. We could modify our approach easily by fixing the number of bits for the result (and intermediates) as well. Either we would allow overflow (computing $(c \times x) \bmod 2^h$), or we would enforce the (intermediate) results to always fit in h bits. However, for SAT we can consider arbitrary bit widths (up to $m_{c \times x}$ bits) for the intermediate results.

3 Applying Single Constant Multiplication to SAT

In this section, we present different approaches for solving the SCM problem. We start with well-known, simple approaches, and then introduce an encoding of SCM into COP. We discuss these approach with a particular focus on their use in SAT encodings.

3.1 A baseline algorithm for encoding $y = c \times x$

In order to establish a baseline, we will first discuss a well-known, simple solution to the SCM problem that does not optimize the resulting circuit at all.

Multiplication by a constant can always be represented using a combination of shifts and additions. Suppose x is defined as a m_x width binary encoded integer. Then we know that $c \times x$ can be at most $c \times (2^{m_x} - 1)$, thus requiring $m_{c \times x} = \lceil \log(c \times (2^{m_x} - 1)) \rceil$ bits to encode.

A basic shift-and-add algorithm will construct a solution to the above SCM problem by decomposing $c \times x$ by a shifted term for every 1 in c 's binary representation:

$$c \times x = \sum_{\substack{0 \leq i < m_c \\ \llbracket \text{bit}(c, i) \rrbracket = 1}} 2^i x$$

The number of additions required is equal to the number of 1s in c 's binary representation minus one, or at most $\lceil \log(c) \rceil - 1$.

3.2 Using Boolean circuit minimization to tackle SCM

Another option for encoding SCM into SAT is to use algorithms such as the Espresso [6] logic minimiser that directly produce small logic circuits based on truth tables or similar representations of Boolean formulae. These algorithms are remarkably powerful, and for small enough problems can produce a guaranteed minimal size circuit for a formula *without introducing auxiliary Boolean variables for intermediate results*.

Example 3. For example, we can construct a CNF encoding for $y = 117 \times x$ with $x \in [0, 15]$ by feeding the following truth table between the x and y bits into Espresso:

	x	y		x	y
0	0000	00000000000000	8	1000	011101010000
1	0001	00000111010101	9	1001	100000011101
2	0010	0001110101010	10	1010	10010010010
3	0011	0010101011111	11	1011	101000000111
4	0100	0011101010100	12	1100	101011111100
5	0101	01001001001001	13	1101	101111110001
6	0110	0101010111110	14	1110	11001100110
7	0111	01100110011	15	1111	11011011011

The circuit minimization of this truth table results in a CNF encoding with 15 variables (simply representing the bits of x and y) and 84 clauses.

Unfortunately as the width w of the input variable x and the size of the constant c grows, the number of Boolean variables in the formula $y = c \times x$ quickly becomes too large for logic minimization. However, for small widths w and constants c logic minimization can generate the best SAT encoding for $y = c \times x$, as demonstrated by our experiments in section 4.

3.3 Formulating SCM as a COP

We will now introduce a general method for solving the SCM problem, by formulating it as a COP.

The following proposition shows that we only need concern ourselves with generating SCM encodings for odd numbers c (since we can derive even multiples by left shifts).

Proposition 1. *Suppose $c = 2^i c'$ where $c' \bmod 2 = 1$ and $i \geq 1$ then we can compute $c \times x$ from $c' \times x$ without any operations.*

Given the above result we can reduce the possible forms of the equation to just three:

Proposition 2. *Given a k equation recipe to compute $c \times x$, $c \bmod 2 = 1$ using the general form above, then there is a k equation recipe using equations of the form $c' = \text{SPLUS}(c_l, s, c_r)$: $c' = 2^s c_l + c_r$, $c' = \text{SMINUS}(c_l, s, c_r)$: $c' = 2^s c_l - c_r$, and $c' = \text{MINUSS}(c_l, s, c_r)$: $c' = c_l - 2^s c_r$.*

Proof. Given a k length list of tuples computing c , we show by induction we can replace this by an equal length list of the new equation forms:

Suppose the first tuple not of these forms is $\langle c_l, s_l, c_r, s_r, \#, c' \rangle$. If $s_r = 0$ then we can rewrite this as either $c' = \text{SPLUS}(c_l, s_l, c_r)$ if $\# = +$ or $c' = \text{SMINUS}(c_l, s_l, c_r)$ if $\# = -$. If $s_l = 0$ then we can rewrite this as either $c' = \text{SPLUS}(c_r, s_r, c_l)$ if $\# = +$ or $c' = \text{MINUSS}(c_l, s_r, c_r)$ if $\# = -$.

Otherwise, $s_l \geq 1 \wedge s_r \geq 1$. Suppose $s_l = s_r$, then we can rewrite the tuple to either $c'' = \text{SPLUS}(c_r, 0, c_l)$ if $\# = +$ or $c'' = \text{SMINUS}(c_l, 0, c_r)$ if $\# = -$, and replace all later usages of c' by c'' by adding s_l to the matching shift argument. Suppose $s_l > s_r$ then we can replace the tuple by recipe $c'' = \text{SPLUS}(c_l, s_l - s_r, c_r)$ if $\# = +$ and $c'' = \text{SMINUS}(c_l, s_l - s_r, c_r)$ if $\# = -$, and replace later uses of c' by c'' adding s_r to the shift argument similarly. Suppose $s_l < s_r$ then we can replace the tuple by recipe $c'' = \text{SPLUS}(c_r, s_r - s_l, c_l)$ if $\# = +$ and $c'' = \text{MINUSS}(c_l, s_r - s_l, c_r)$ if $\# = -$, and replace later uses of c' by c'' adding s_r to the shift argument similarly.

The requirement that $c \bmod 2 = 1$ means that one of the first two cases is applicable to the last equation. $\square$

Example 4. Both recipes for 117 of example 2 can be represented equivalently as $[17 = \text{SPLUS}(1, 4, 1), 19 = \text{SPLUS}(1, 1, 17), 117 = \text{SMINUS}(17, 3, 19)]$. The second recipe is defined by $[63 = \text{SMINUS}(1, 6, 1), 59 = \text{MINUSS}(63, 2, 1), 117 = \text{SMINUS}(59, 1, 1)]$.

A MiniZinc [13] model for solving SCM as a combinatorial optimization problem is shown in listing 1. Note that it only makes use of the three recipes from proposition 2. Each equation is represented by its type **ty**, with **NOP** added for an unused equation; the equation numbers defining its left (**left**) and right (**right**) arguments, and the shift amount **shift**.⁵ We use equation number 0 to represent $1 \times x$. The multiplier resulting from an equation is defined by **mult**. Line 14 sets the dummy equation multiplier to 1, and enforces the final result to be the target. In lines 16-18 we enforce that all equations after **used** are NOPs, and set their other components to dummy values. Line 19 enforces that equations use only earlier defined multiples. The computation of the multiplier for each equation (lines 20-27) follows the definition, where **NOP** just returns the previous multiplier. We minimize the number of used equations. Note that the model “pre-computes” the powers of 2 in **p2** (line 13), since solvers typically propagate poorly for exponential expressions.

3.4 Encoding an SCM recipe

Given a solution (e.g., example 2) to the SCM problem, we can encode the shifts and adds of each equation ($c' \times x = 2^{s_l} \times (c_l \times x) \pm 2^{s_r} \times (c_r \times x)$). We obtain the

⁵ The careful reader will note that we limit the maximum shift using the target. Relaxing this restriction and never found better recipes (in terms of equations or adders), but we have no formal proof for it.

```

1 int: xbits; % number of bits to represent x
2 int: n;      % max number of shift +- equations
3 set of int: EQ = 1..n;
4 set of int: EQ0 = 0..n; % equation number 0 = 1x
5 int: maxsh = ceil(log2(target)); % maximum left shift
6 int: target; % target multiplier
7 enum TYPE = { SPLUS, SMINUS, MINUSS, NOP };
8 array[EQ] of var TYPE: ty; % type of equation
9 array[EQ] of var EQ0: left; % left input equation number
10 array[EQ] of var EQ0: right; % right input equation number
11 array[EQ] of var 0..maxsh: shift; % shift left applied
12 array[EQ0] of var 0..infinity: mult; % multiplier value
13 array[1..maxsh+xbits] of int: p2 = [2^i | i in 1..maxsh+xbits];
14 constraint mult[0] = 1 /\ mult[n] = target;
15 var EQ0: used; % number of equations used
16 constraint forall(e in EQ)(e > used <-> ty[e] = NOP);
17 constraint forall(e in EQ)(e > used ->
18     left[e] = 0 /\ right[e] = 0 /\ shift[e] = 1);
19 constraint forall(e in EQ)(left[e] < e /\ right[e] < e);
20 constraint forall(e in EQ)(mult[e] =
21     if ty[e] = SPLUS then
22         p2[shift[e]] * mult[left[e]] + mult[right[e]]
23     elseif ty[e] = SMINUS then
24         p2[shift[e]] * mult[left[e]] - mult[right[e]]
25     elseif ty[e] = MINUSS then
26         mult[left[e]] - p2[shift[e]] * mult[right[e]]
27     else mult[e-1] endif);
28 solve minimize used; % minimize equations

```

Listing 1: A MiniZinc model to find a recipe of at most size n for computing $\text{target} \times x$ where x is represented by xbits bits.

binary encoding of the output $z = (c' \times x)$ from the already computed binary encodings of the inputs $z_l = (c_l \times x)$ and $z_r = (c_r \times x)$.

To apply a shift $y_l = z_l \ll s$ (or equivalently, $y_l = 2^s z_l$) on the binary encoding of z_l (or to apply $y_r = z_r \ll s$), we simply extend its encoding by $\llbracket \text{bit}(y_l, i) \rrbracket = 0, 0 \leq i < s$ and $\llbracket \text{bit}(y_l, i) \rrbracket = \llbracket \text{bit}(z_l, i - s) \rrbracket, s \leq i \leq m_{z_l}$ using no additional variables or clauses.

After shifting both inputs, it remains to encode an addition of the form $z = y_l + y_r$. Here, we can encode a ripple carry adder [17] which produces the output (sum) bits $\llbracket \text{bit}(z, i) \rrbracket$, and auxiliary carry bits $\llbracket \text{bit}(c, i) \rrbracket$. Inputs variables y_l, y_r might contain fixed literals due to the applied shifts, which in turn leads to output variables that are fixed, or that are equivalent to input bits, or their negations. The initial carry bit is also fixed as $\llbracket \text{bit}(r, 0) \rrbracket = 0$ (which is essentially why a half-adder is used to add $\llbracket \text{bit}(y_l, 0) \rrbracket$ and $\llbracket \text{bit}(y_r, 0) \rrbracket$). We optimize the ripple carry adder encoding to account for this.

The current sum bit $\llbracket \text{bit}(z, i) \rrbracket$ is the result of the **xor** operation on up to three input variables, $\llbracket \text{bit}(y_l, i) \rrbracket \oplus \llbracket \text{bit}(y_r, i) \rrbracket \oplus \llbracket \text{bit}(r, i) \rrbracket$, which we reformulate more generically as $\oplus(\llbracket \text{bit}(y_l, i) \rrbracket, \llbracket \text{bit}(y_r, i) \rrbracket, \llbracket \text{bit}(r, i) \rrbracket)$.

We remove any fixed variables from the **xor**'s input, and compute the sum of their values. Then, we apply the **xor** on the remaining non-fixed variables, negating the output if the fixed sum is odd. Given a list l of bits, some of which are fixed, define $\text{fs}(l)$ as the sum of the fixed bits in l and $\text{vb}(l)$ the list of unfixed (variable) bits in l . We can encode the $\llbracket \text{bit}(z, i) \rrbracket = \oplus(\llbracket \text{bit}(y_l, i) \rrbracket, \llbracket \text{bit}(y_r, i) \rrbracket, \llbracket \text{bit}(r, i) \rrbracket)$ as follows:

$$\oplus(l) = \begin{cases} \text{fs}(l) \bmod 2, & \text{if } \text{vb}(l) = [] \\ b_1, & \text{if } \text{vb}(l) = [b_1] \wedge \text{fs}(l) = 0 \\ \neg b_1, & \text{if } \text{vb}(l) = [b_1] \wedge \text{fs}(l) = 1 \\ b_1 \oplus b_2 & \text{if } \text{vb}(l) = [b_1, b_2] \wedge \text{fs}(l) = 0 \\ b_1 \oplus b_2 \oplus 1 & \text{if } \text{vb}(l) = [b_1, b_2] \wedge \text{fs}(l) = 1 \\ b_1 \oplus b_2 \oplus b_3 & \text{otherwise } \text{vb}(l) = [b_1, b_2, b_3] \end{cases}$$

If $\oplus(l)$ is constant or a literal, we can just equate $\llbracket \text{bit}(z, i) \rrbracket$ with the result, requiring no encoding clauses. Otherwise, $\llbracket \text{bit}(z, i) \rrbracket$ is equated to a $\oplus$ expression, which can be encoded using a single XOR clause if supported by the SAT solver [16], or encoded with (standard) clauses in a well understood manner.

The next carry bit is set to true if at least 2 of its input variables are. In other words, $\llbracket \text{bit}(r, i+1) \rrbracket$ is true if and only if the expression $\llbracket \text{bit}(y_l, i) \rrbracket + \llbracket \text{bit}(y_r, i) \rrbracket + \llbracket \text{bit}(r, i) \rrbracket \geq 2$ is true. Again, we split off and sum the fixed variables, and compute the carry by evaluating the generic version of the at-least-2 constraint as follows:

$$\left(\sum(l) \geq 2 \right) = \begin{cases} 1, & \text{if } \text{fs}(l) \geq 2 \\ 0 & \text{if } \text{vb}(l) = [] \wedge \text{fs}(l) < 2 \\ 0, & \text{if } \text{vb}(l) = [b_1] \wedge \text{fs}(l) = 0 \\ b_1, & \text{if } \text{vb}(l) = [b_1] \wedge \text{fs}(l) = 1 \\ b_1 \wedge b_2, & \text{if } \text{vb}(l) = [b_1, b_2] \wedge \text{fs}(l) = 0 \\ b_1 \vee b_2, & \text{if } \text{vb}(l) = [b_1, b_2] \wedge \text{fs}(l) = 1 \\ (b_1 \wedge b_2) \vee (b_1 \wedge b_3) \vee (b_2 \wedge b_3) & \text{otherwise } \text{vb}(l) = [b_1, b_2, b_3] \end{cases}$$

If $\sum(l) \geq 2$ is constant or a literal, we can just equate $\llbracket \text{bit}(r, i+1) \rrbracket$ with the result, requiring no encoding clauses. Otherwise, $\llbracket \text{bit}(r, i+1) \rrbracket$ is equated to the CNF expression as normal.

To handle subtraction, we use the fact that ripple-carry adders directly handle 2's complement. So, $y_l + (-y_r) = y_l + (\bar{y}_r + 1)$ where the encoding of $\bar{y}_r$ is the complement of y_r , i.e., $\llbracket \text{bit}(\bar{y}_r, i) \rrbracket = \neg \llbracket \text{bit}(y_r, i) \rrbracket$. To offset the constant 1, we can add an initial carry of 1.

3.5 Minimizing the number of adders

The COP formulation of section 3.3 minimizes the number of equations, k . However, we show in section 3.4 how the true encoding complexity for each equation is more closely related to the number of (full/half) adders. In this section, we formulate an alternative objective which minimizes the number of adders a instead.

We are now required to consider the bit width w of the input x being multiplied as an additional parameter. Consider the addition $c' \times x = 2^{s_l} \times (c_l \times x) + 2^{s_r} \times (c_r \times x)$. Assuming the left input $c_l \times x$ is represented in $m_{c_l \times x}$ bits and the right input $c_r \times x$ is represented in $m_{c_r \times x}$ bits, then the number of full/half adders naïvely required for the addition is the width of the result $m_{c' \times x}$. However, we can omit full/half adders for low bits where one argument is guaranteed 0, and if the non-zero bits never overlap we need no adders. So the number of full/half adders required is 0 if $s_l > m_{c_r \times x} + s_r$ or if $s_r > m_{c_l \times x} + s_l$, otherwise the number of adders required is $\max(m_{c_l \times x} + s_l, m_{c_r \times x} + s_r) - \max(s_l, s_r)$. For subtractions, the only bits not requiring adders are the low bits where both operands are known to be 0, so we require $m_{c' \times x} - \min(s_l, s_r)$ full/half adders.

We can similarly show that we do not need the more general recipes, we can restrict to the recipes of proposition 2, since these also do not change the number of full/half adders required.

Proposition 3. *Given an a adder recipe to compute $c \times x, c \bmod 2 = 1$ using the general form of k equations, then there is an a adder recipe using equations of the form $c' = \text{SPLUS}(c_l, s, c_r): c' = 2^s c_l + c_r$, $c' = \text{SMINUS}(c_l, s, c_r): c' = 2^s c_l - c_r$, and $c' = \text{MINUSS}(c_l, s, c_r): c' = c_l - 2^s c_r$.*

Proof. (Sketch) The proof follows that of proposition 2. Note that if we shift both addends right by the same amount, then we don't change the number of adders required since the low bits do not require adders. When we replace the addition of c' by c'' shifted left by some amount s , we are simply computing the same addition since $c' = 2^s c''$ which does not change the number of adders required. $\square$

We can modify our MiniZinc model (see listing 1) to keep track of the number of bits required for the result of each equation, and count the number of full/half adders used for each equation, and make that the new objective. The additions to the model (and replacement objective) are shown in listing 2. The number of bits required for the result of each equation are defined by `bits` while the number of full/half adders for each equation are given by `adders`. In lines 4-5 we compute the bit width directly from the multiple (the least power big enough to hold the maximum result). We compute the number of adders required for each recipe in lines 6-12. Finally, we minimize the sum of adders required.

Note that for minimizing the SAT model we only minimize a proxy for the size of the resulting SAT model, the number of full/half adders. We experimented with directly minimizing the number of clauses or literals (sum of size of clauses) in the SAT encoding. The resulting models were much harder to solve, we usually couldn't prove optimality, and for the small examples where we could the result

```

1 array[EQ] of var 0..infinity: bits; % number of bits
2 array[EQ] of var 0..infinity: adders; % number of adders
3 include "arg_max.mzn";
4 constraint forall(e in EQ)(bits[e] =
5   arg_max([mult[e]*(p2[xbits]-1)<=p2[i]|i in 1..maxsh+xbits ]));
6 constraint forall(e in EQ)(adders[e] =
7   if ty[e] = SPLUS then
8     if shift[e] >= bits[right[e]] then 0
9     else max(bits[left[e]],bits[right[e]]-shift[e]) endif
10  elseif ty[e] = SMINUS then bits[e]
11  elseif ty[e] = MINUSS then bits[e]
12  else 0 endif);
13 solve minimize sum(adders); % minimize adders

```

Listing 2: Additions to the MiniZinc model of listing 1 to optimize for the number of full/half adders.

was rarely better than the SAT encoding resulting from minimizing full/half adders. It remains interesting future work to see if this more direct approach to encoding minimization could be improved.

4 Experimental evaluation

To evaluate the effectiveness of the proposed approach, we compare the size of the encoding of $c \times x$ for a range of constants c and bit widths for x in section 4.1. In section 4.2 we apply our approach to a realistic benchmark, comparing both the encoding size and solve times.

4.1 Constructing and analyzing the SCM databases

In this section, we compare the different SCM approaches in terms of the size of their respective encodings. For each bit width $2 \leq w \leq 16$, we average the number of variables, clauses, and literals that are generated when encoding $c \times x$ over $1 \leq c \leq 2047$. The four methods are **base** (the baseline from section 3.1), **espresso** (using logic minimization with Espresso 2.3 from section 3.2), **min-k** (minimizing the number of ripple carry adders) and **min-a** (minimizing the total number of adders). It becomes increasingly hard for Espresso to compute CNF for all constants as the number of input bits grows, so for **espresso** we are limited to $w \leq 12$. The results are shown in fig. 1.

It is clear that the naive decomposition of **base** produces by far the most variables. In contrast, **min-k** achieves the same multiplication using fewer ripple carry adders. Importantly, **min-a** has even fewer variables than **min-k**, indicating that minimizing the number of adders serves as a proxy for avoiding additional variables. The effect is stronger for smaller bit widths, presumably because then there is more opportunity for completely free additions where one input is shifted

by the other input’s bit width. The number of variables for **espresso** is always minimal, since no additional variables are ever constructed. In terms of clauses and literals, we see a linear relationship between the other methods other than **espresso**, which generates an exponential number of clauses and literals, **min-a** better than **min-k** better than **base**.




Fig. 1: Encoding size of $c \times x$ using different SCM approaches

4.2 Solving Multidimensional Bounded Knapsack Problem

To evaluate the different approaches, we adapt the approach from [10] to generate hard *Multidimensional Bounded Knapsack Problem* (MBKP) instances. In MBKP, we decide for N item types how many x_i to pack (up to B) such that a minimum profit $\sum_{i=1}^N x_i p_i \geq P$ is reached. Each item is restricted by M dimensions of weight with $\bigwedge_{i=1}^M \sum_{j=1}^N x_j w_{i,j} \leq W_i$. Instances can be generated by first generating C coefficient sets for $w_{i,j}$ and p_i , sampling uniformly from $[1, Q]$. Then, for each coefficient set, we choose S capacity sets for $1 \leq s \leq S$ with a capacity factor $f = \frac{s}{S}$: generating $W_i = f \sum_{j=1}^N B w_{i,j}$, $1 \leq i \leq M$, and $P = (1 - f) \sum_{j=1}^N B p_j$. As f increases, the constraints become less strict, and instances turn from unsatisfiable to satisfiable. To avoid an excess of trivial instances on the lower and higher ends of f , we normalize f to be within $0.2 \leq f \leq 0.8$. One instance set is generated for each $B = 2^w - 1$, $4 \leq w \leq 16$, with parameters $N = 15$, $M = 100$, $S = 100$, $C = 3$.

In the SAT encoding of MBKP each integer variable x_i is binary encoded in the least number of bits required to represent its full domain. Each linear constraint of the form $\sum_{j=1}^N w_{i,j} x_j \leq W_i$ is encoded by first constructing $y_j = w_{i,j} \times x_j$ using the SCM method, and then computing $y_0 = \sum_{j=1}^N y_j$ by repeated use of ripple-carry adders, and finally constraining $y_0 \leq W_i$ using a lexicographic constraint.

To provide another control method, we encode the benchmarks using Picat-SAT (version 3.5) [18], a state-of-the-art SAT encoder that uses the binary encoding. The **picat** encodes SCM similar to **base**, but with additional ad-hoc

optimizations [19]. Note that for these problems with large coefficients and domain sizes alternate encoding approaches, e.g., domain or order encodings are non-competitive.

All encodings produced in our experiments are solved using the same binary of the SAT solver CaDiCal (version 1.9.1) [4] with default parameters, 8 GB memory limit and a time limit of 180 seconds. A PAR2 penalty applies to time-outs.

In fig. 2, we compare the encoding sizes of the MBKP, similar to the results from section 4.1. **picat** matches **min-a** in the number of variables. However, for clauses and especially for literals, larger bit widths result in larger encodings, closer to **base**. For unknown reasons, **picat** shows a temporary spike in clauses and literals at $w = 5$ (for some constraints, **picat** also relies on Espresso; which it perhaps uses in this case).




Fig. 2: Encoding size of MBKP using different SCM approaches

Comparing solve times in fig. 3, we observe that **espresso** performs best up to a bit width of 10, after which its performance drops off drastically. Even if pre-computing larger input bit widths would be tractable for **espresso**, we can extrapolate that an encoding without auxiliary variables loses its effectiveness. In contrast, **min-k** and **min-a** remain relatively constant, even though the instances grow exponentially in size. Furthermore, **min-a** performs better than **min-k**, which shows the benefit of minimizing adders over ripple carry adders, even if some recipes have sub optimal length (such as in example 1). Finally, **picat** performs much better than **base** for smaller bit widths but similar for larger bit widths, suggesting that **picat** is able to apply certain encoding optimizations in the former case.

5 Conclusion and Future Work

In conclusion, we have shown how to tackle and apply SCM specifically for SAT in order to encode linear constraints. Encoding linear constraints is a key challenge for SAT-based solvers [1]. To our knowledge, this is the first work that adapts SCM circuits to encode linear constraints for SAT. Since ripple carry




Fig. 3: Solve time comparison for MBKP

adders are not uniformly encoded in the presence of partially fixed inputs, this gives us the opportunity to develop different optimal circuits compared to the traditional SCM problem. In the experimental evaluation, we have seen how this approach significantly improves both the encoding size as well as the solver’s performance. Finally, we have made the pre-computed SCM encodings freely available.

In future work, we aim to extend our approach by applying MCM to SAT, which – unlike SCM – cannot be comprehensively pre-computed. Instead, we would first have to collect the target constants (e.g., through common sub-expression elimination [14]) from a given instance and then dynamically solve the MCM problem at encode time. This poses an interesting trade-off between encode and solve time, but has the potential for additional improvements.

References

1. Abío, I., Mayer-Eichberger, V., Stuckey, P.J.: Encoding Linear Constraints into SAT. CoRR **abs/2005.02073** (2020), <https://arxiv.org/abs/2005.02073>
2. Aksoy, L., Flores, P.F., Monteiro, J.: Exact and Approximate Algorithms For The Filter Design Optimization Problem. IEEE Trans. Signal Process. **63**(1), 142–154 (2015). <https://doi.org/10.1109/TSP.2014.2366713>, <https://doi.org/10.1109/TSP.2014.2366713>
3. Avizienis, A.: Signed-Digit Number Representations For Fast Parallel Arithmetic. IRE Transactions on Electronic Computers **EC-10**(3), 389–400 (1961). <https://doi.org/10.1109/TEC.1961.5219227>
4. Biere, A., Fazekas, K., Fleury, M., Heisinger, M.: CaDiCaL, Kissat, Paracooba, Plingeling and Treengeling Entering the SAT Competition 2020. In: Balyo, T., Froleys, N., Heule, M., Iser, M., Jarvisalo, M., Suda, M. (eds.) Proc. of SAT Competition 2020 – Solver and Benchmark Descriptions. Department of Computer Science Report Series B, vol. B-2020-1, pp. 51–53. University of Helsinki (2020)

5. Biere, A., Heule, M., van Maaren, H., Walsh, T. (eds.): Handbook of Satisfiability. IOS Press, 2 edn. (May 2021), google-Books-ID: dUAvEAAAQBAJ
6. Brayton, R.K., Hachtel, G.D., McMullen, C.T., Sangiovanni-Vincentelli, A.L.: Logic Minimization Algorithms for VLSI Synthesis, The Kluwer International Series in Engineering and Computer Science, vol. 2. Springer (1984). <https://doi.org/10.1007/978-1-4613-2821-6>, <https://doi.org/10.1007/978-1-4613-2821-6>
7. Cappello, P., Steiglitz, K.: Some Complexity Issues in Digital Signal Processing. IEEE Transactions on Acoustics, Speech, and Signal Processing **32**(5), 1037–1041 (1984)
8. Dekker, J.J., Bierlee, H.: Pindakaas: CPAIOR-24 (Mar 2024). <https://doi.org/10.5281/zenodo.10851856>, <https://doi.org/10.5281/zenodo.10851856>
9. Gustafsson, O.: Towards Optimal Multiple Constant Multiplication: A Hypergraph Approach. In: 42nd Asilomar Conference on Signals, Systems and Computers, ACSSC 2008, Pacific Grove, CA, USA, October 26-29, 2008. pp. 1805–1809. IEEE (2008). <https://doi.org/10.1109/ACSSC.2008.5074738>, <https://doi.org/10.1109/ACSSC.2008.5074738>
10. Han, B., Leblet, J., Simon, G.: Hard Multidimensional Multiple choice Knapsack Problems, an Empirical Study. Comput. Oper. Res. **37**(1), 172–181 (2010). <https://doi.org/10.1016/j.cor.2009.04.006>, <https://doi.org/10.1016/j.cor.2009.04.006>
11. Kumm, M.: Optimal Constant Multiplication Using Integer Linear Programming. IEEE Trans. Circuits Syst. II Express Briefs **65-II**(5), 567–571 (2018). <https://doi.org/10.1109/TCSII.2018.2823780>, <https://doi.org/10.1109/TCSII.2018.2823780>
12. Ma, S., Ampadu, P.: Optimal SAT-based Minimum Adder Synthesis of Linear Transformations. In: Lee, H., Geiger, R.L. (eds.) 62nd IEEE International Midwest Symposium on Circuits and Systems, MWSCAS 2019, Dallas, TX, USA, August 4-7, 2019. pp. 335–338. IEEE (2019). <https://doi.org/10.1109/MWSCAS.2019.8885033>, <https://doi.org/10.1109/MWSCAS.2019.8885033>
13. Nethercote, N., Stuckey, P., Becket, R., Brand, S., Duck, G., Tack, G.: MiniZinc: Towards a Standard CP Modelling Language. In: Bessiere, C. (ed.) Proceedings of the 13th International Conference on Principles and Practice of Constraint Programming. LNCS, vol. 4741, pp. 529–543. Springer-Verlag (2007)
14. Nightingale, P., Spracklen, P., Miguel, I.: Automatically Improving SAT Encoding of Constraint Problems Through Common Subexpression Elimination in Savile Row. In: Pesant, G. (ed.) Principles and Practice of Constraint Programming - 21st International Conference, CP 2015, Cork, Ireland, August 31 - September 4, 2015, Proceedings. Lecture Notes in Computer Science, vol. 9255, pp. 330–340. Springer (2015). https://doi.org/10.1007/978-3-319-23219-5_23, https://doi.org/10.1007/978-3-319-23219-5_23
15. Rossi, F., van Beek, P., Walsh, T. (eds.): Handbook of Constraint Programming, Foundations of Artificial Intelligence, vol. 2. Elsevier (2006), <https://www.sciencedirect.com/science/bookseries/15746526/2>
16. Soos, M., Nohl, K., Castelluccia, C.: Extending SAT Solvers to Cryptographic Problems. In: Kullmann, O. (ed.) Theory and Applications of Satisfiability Testing - SAT 2009, 12th International Conference, SAT 2009, Swansea, UK, June 30 - July 3, 2009. Proceedings. Lecture Notes in Computer Science, vol. 5584, pp.

- 244–257. Springer (2009). https://doi.org/10.1007/978-3-642-02777-2_24, https://doi.org/10.1007/978-3-642-02777-2_24
17. Warners, J.P.: A Linear-time Transformation of Linear Inequalities into Conjunctive Normal Form. *Inf. Process. Lett.* **68**(2), 63–69 (1998). [https://doi.org/10.1016/S0020-0190\(98\)00144-6](https://doi.org/10.1016/S0020-0190(98)00144-6), [https://doi.org/10.1016/S0020-0190\(98\)00144-6](https://doi.org/10.1016/S0020-0190(98)00144-6)
 18. Zhou, N., Kjellerstrand, H.: The Picat-SAT Compiler. In: Gavanelli, M., Reppy, J.H. (eds.) *Practical Aspects of Declarative Languages - 18th International Symposium, PADL 2016, St. Petersburg, FL, USA, January 18-19, 2016. Proceedings. Lecture Notes in Computer Science*, vol. 9585, pp. 48–62. Springer (2016). https://doi.org/10.1007/978-3-319-28228-2_4, https://doi.org/10.1007/978-3-319-28228-2_4
 19. Zhou, N., Kjellerstrand, H.: Optimizing SAT Encodings for Arithmetic Constraints. In: Beck, J.C. (ed.) *Principles and Practice of Constraint Programming - 23rd International Conference, CP 2017, Melbourne, VIC, Australia, August 28 - September 1, 2017, Proceedings. Lecture Notes in Computer Science*, vol. 10416, pp. 671–686. Springer (2017). https://doi.org/10.1007/978-3-319-66158-2_43, https://doi.org/10.1007/978-3-319-66158-2_43